

PATENT
EXPRESS MAIL NO. EV328618202US

ASSOCIATIVE DATABASE SCANNING AND INFORMATION RETRIEVAL

Background of the Invention

Associative memory devices are known in the prior art. Generally, these associative memory devices comprise peripheral memories for computers, computer networks, and the like, which 5 operate asynchronously to the computer, network, etc. and provide increased efficiency for specialized searches. Additionally, it is also known in the prior art that these memory devices can include certain limited decision-making logic as an aid to a main CPU in accessing the peripheral memory. An example of such an associative 10 memory device particularly adapted for use with a rotating memory such as a high speed disk or drum can be found in U.S. Patent No. 3,906,455, the disclosure of which is incorporated herein by reference. This particular device provides a scheme for use with a rotating memory and teaches that two passes over a memory sector is 15 necessary to presort and then sort the memory prior to performing any logical operations thereon. Thus, this device is taught as not being suitable for use with any linear or serial memory such as magnetic tape or the like.

Other examples of prior art devices may also be found in U.S. 20 Patent Nos. 3,729,712; 4,464,718; 5,050,075; 5,140,692; and 5,721,898; the disclosures of which are incorporated herein by reference.

As an example, in 4,464,718, Dixon performs fixed comparisons on a fixed number of bytes. They don't have the ability to scan and correlate arbitrarily over the data. They search serially along the tracks in a given disk cylinder but there is no provision for 5 parallel searching across disks. Dixon's comparisons are limited by a fixed rigid number of standard logical operation types. Additionally, the circuitry presented supports only these single logical operations. There is no support for approximate or fuzzy matching.

10 While these prior art associative memory devices represent an attempt to speed the input and output of information to and from a peripheral memory, which in many cases is a mass storage memory device, all rely on the classic accessing of data stored in digital form by reading and interpreting the digital either address or 15 content of the memory location. In other words, most such devices access data by its address but there are some devices that take advantage of the power of content addressing as is well known in the art. Nevertheless, in all of the prior art known to the inventors, the digital value of the address or data contained in the addressed 20 location must be read and interpreted in its digital form in order to identify the data and then select it for processing. Not only does it take processing time to read and interpret the digital data represented by the address or content, this necessarily requires that the accessing circuit process the memory according to the structure 25 of the data stored. In other words, if the data is stored in octets, then the accessing circuitry must access the data in octets and process it in an incremental manner. This "start and stop" processing serves to increase the input/output time required to access data. As is also well known in the art, this input/output 30 time typically represents the bottleneck and effective limitation of processing power in any computer or computer network.

In order to solve these and other problems in the prior art, the inventors herein have succeeded in designing and developing a method and apparatus for an associative memory in several embodiments 35 which provide an elegantly simple solution to these prior art limitations as well as dramatically decreased access times for data stored in mass storage memories. As mentioned, the invention has several embodiments each of which has its own advantages.

Generally, the invention may be described as a technique for 40 data retrieval through approximate matching of a data key with a continuous reading of data as stored on a mass storage medium. In

its broadest, and perhaps most powerful, embodiment, this key may be an analog signal and it is matched with an analog signal generated by a typical read/write device as it slews across the mass storage medium. In other words, the steps taught to be required in the prior art of not only reading the analog representation of digital data stored on the mass storage medium but also the conversion of that signal to its digital format prior to being compared are eliminated. Furthermore, there is no requirement that the data be "framed" or compared utilizing the structure or format in which the data has been organized and stored. For an analog signal, all that need be specified is the elapsed time of that signal which is used for comparison with a corresponding and continuously changing selected time portion of the "read" signal. Using any one of many standard correlation techniques as known in the prior art, the data "key" may then be approximately matched to the sliding "window" of data signal to determine a match. Significantly, the same amount of data may be scanned much more quickly and data matching the search request may be determined much more quickly as well. For example, the inventors have found that CPU based approximate searches of 200 megabytes of DNA sequences can take up to 10 seconds on a typical present day "high end" system, assuming the offline processing to index the database has already been completed. In that same 10 seconds, the inventors have found that a 10-gigabyte disk could be magnetically searched for approximate matches using the present invention. This represents a 50:1 improvement in performance. Furthermore, in a typical hard disk drive there are four surfaces and corresponding read/write heads, which may be all searched in parallel should each head be equipped with the present invention. As these searches can proceed in parallel, the total increase in speed or improvement represents a 200:1 advantage. Furthermore, additional hard disk drives may be accessed in parallel and scaled to further increase the advantage provided by the present invention.

By choosing an appropriate correlation or matching technique, and by setting an appropriate threshold, the search may be conducted to exactly match the desired signal, or more importantly and perhaps more powerfully, the threshold may be lowered to provide for approximate matching searches. This is generally considered a more powerful search mode in that databases may be scanned to find "hits" which may be valid even though the data may be only approximately that which is being sought. This allows searching to find data that has been corrupted, incorrectly entered data, data which only

generally corresponds to a category, as well as other kinds of data searches that are highly desired in many applications. For example, a library of DNA sequences may be desired to be searched and hits found which represent an approximate match to a desired sequence of 5 residues. This ensures that sequences which are close to the desired sequence are found and not discarded but for the difference in a forgivable number of residue mismatches. Still another application involves Internet searches provided by Internet search engines. In such a search, approximate matching allows for misspelled words, 10 differently spelled words, and other variations to be accommodated without defeating a search or requiring a combinatorial number of specialized searches. This technique permits a search engine to provide a greater number of hits for any given search and ensure that a greater number of relevant web pages are found and cataloged in the 15 search.

Still another possible application for the present invention is for accessing databases which may be enormous in size or which may be stored as analog representations. For example, our society has seen the implementation of sound recording devices and their use in many 20 forums including judicial proceedings. In recent history, tape recordings made in the President's oval office have risen in importance with respect to impeachment hearings. As can be appreciated, tape recordings made over the years of a presidency can accumulate into a huge database which might require a number of 25 persons to actually listen to them in order to find instances where particular words are spoken that might be of interest. Utilizing the present invention, an analog representation of that spoken word can be used as a key and sought to be matched while the database is scanned in a continuous manner and at rapid speed. Thus, the present 30 invention provides a powerful search tool for massive analog databases as well as massive digital databases.

While text-based searches are accommodated by the present invention as described above, storage media containing pictures, sound, and other representations have traditionally been more 35 difficult to search than text. The present invention allows searching a large data space for the presence of such content or fragments thereof. For example, the key in this case could be a row or quadrant of pixels that represent the image being sought. Approximate matching of the key's signal can then allow 40 identification of matches or near matches to the key.

While the principal advantages and features of the present invention have been briefly explained above, a more thorough understanding of the invention may be attained by referring to the drawings and description of the preferred embodiment which follow.

5

Brief Description of the Drawings

Figure 1 is a block diagram of a computer having access through its system or input/output bus to the present invention comprising the logical mass storage medium for asynchronous access thereto, as 10 well as a network interface between the present invention and a network to other computer systems;

Figure 1A is a schematic for a conventional rigid disk drive system illustrating different insertion points for connection of the present invention;

15 Figure 2 is a block diagram depicting the implementation of the present invention in a stand-alone configuration;

Figure 3 is a block diagram depicting the present invention implemented as a shared remote mass storage device across a network;

20 Figure 4 is a block diagram depicting the present invention as a network attached storage device (NASD);

Figure 5 is a flow chart detailing the logical steps in the inventive method for accessing data in a mass storage medium;

Figure 6 is a graphical representation of an analog signal as might be used as a key;

25 Figure 7 is a graphical representation of an analog signal representing the continuous reading of data from a mass storage medium in which the "key" data pattern is present;

Figure 8 is a graphical representation of the key signal overlying and matched to the data signal;

30 Figure 9 is a graphical representation of a correlation function calculated continuously as the data in the mass storage medium is scanned and compared with the key; and

Figure 10 is a graphical representation of a correlation function as the key is continuously compared with a signal taken from 35 reading a different set of data but which also contains the "key" data pattern.

Detailed Description of the Preferred Embodiment

As shown in Figure 1, the present invention is readily 40 implemented in a stand-alone computer or computer system. In broad terms, the present invention is comprised of an approximate matching

and pre-fetch processor 20 which itself comprises a programmable digital logic device or microprocessor 22 coupled to an approximate matching unit 24. A mass storage medium 26 is connected to the approximate matching unit 24 and provides the medium for storing 5 large amounts of data. The term "mass storage medium" should be understood as meaning any device used to store large amounts of data, and which is typically designated for use in a computer or computer network. Examples include optical systems such as CD ROMS, magnetic systems such as hard disk drives or sub-units such as a single disk 10 surface, and these systems may be rotating, linear, serial, parallel, or various combinations of each. For example, a rack of hard disk drive units could be connected in parallel and their parallel output provided at the transducer level to one or more approximate matching units 24. Similarly, a bank of magnetic tape drives could be used, 15 and their serial outputs each provided in parallel to one or more approximate matching units 24. The data stored on the medium may be in analog or in digital form. For example, the data could be voice recordings as in the tape recording example given above. The present invention is thus scalable, permitting an increase in the amount of 20 data stored in a mass storage system (by increasing the number of parallel storage units) while preserving the performance (by increasing the number of parallel approximate matching units or replicating the approximate matching and prefetch processor).

In the prior art as shown in the upper portion of Figure 1, 25 typically a disk controller 28 and/or a disk cache 30 may be used in the traditional sense for access by a CPU 32 over its system or input/output bus 34. In the present invention, the approximate matching and pre-fetch processor 20 accesses data in the mass storage medium 26 and presents it for use at the system bus 34 without moving 30 large blocks of memory from the mass storage medium 26 over the system bus 34 and into the main memory (not separately shown) of CPU 32 for sorting and accessing. In other words, as is explained in greater detail below, the CPU 32 sends a data request or query to the approximate matching and pre-fetch processor 20 which then 35 asynchronously accesses and sorts data in the mass storage medium 26 and presents it for use either in a disk cache 30 as is known in the prior art or directly onto the system bus 34 without further processing being required by CPU 32 or use of its main memory. The processor is thus free to perform other tasks while the searching and 40 matching activity is being performed by the present invention.

The approximate matching unit 24 may itself comprise a set of digital logic or analog circuitry organized to perform scanning, correlation, general logic and approximate matching functions on the signals received from the storage medium in conjunction with digital or analog representations of query directed key information. The functions may be performed by dedicated logic and circuitry, programmable logic and circuitry or by use of programmable processors such as digital signal processors (DSPs). The inventors contemplate that many different variations of hardware or programmed software devices may be used to virtually equal advantage in implementing the present invention and therefore do not contemplate any particular implementation as a better mode than another, at present, except as otherwise expressed herein. Furthermore, the present invention should not be considered as being tied to any specific data structure or application, those having been used herein are instead to be viewed as illustrative only. Using the present disclosure as a template, it is contemplated that one of ordinary skill in the art could readily construct a device as would perform the functions and operation necessary to achieve the purposes of the invention.

As has been explained above, the present invention may be used to perform both exact matches and approximate matches. When performing an exact match in the analog domain, at Point A in Figure 1A, where matching is done using analog comparators and correlation techniques, there an exact match corresponds to setting a sufficiently high threshold value for matching the key with analog source data on the storage medium. Approximate matching in this analog domain corresponds to setting appropriate (lesser) threshold values. When performing an exact match in the digital domain, an exact match is performed using digital comparators and logic as suggested in current art, where a digital key is compared with digital source data from the storage medium. Such matching could be performed at Point B or Point C, as shown in Figure 1A, which corresponds to the pre- and post-error-corrected digital signal, respectively. Approximate matching in this domain corresponds to performing comparisons or digital correlation of the digital key with digital source data obtained from the storage medium. The digital key may contain "wild card" or "don't care" digital bits or data formats. The success of an approximate match may be determined by setting a correlation value or by using one of a number of matching-performance metrics such as the number of bits within a key that are equal to the corresponding bits in the scanned source data. Also,

note that the data entries identified in an "approximate" match search will include the "exact" hits that would result from an "exact" search. For clarity, when the word "match" is used, the reader should understand that it includes a search or a data result found through either of an approximate search or an exact search. When the phrase "approximate match" or even just "approximate" is used, the reader should understand that it could be either of the two searches described above as approximate searches, or for that matter any other kind of "fuzzy" search that has a big enough net to gather data entries that are loosely related to the search criteria. Of course, an exact match is just that, and does not include any result other than an exact match of the search criteria with a high degree of correlation.

The microprocessor or programmable logic device 22 may itself comprise one or more microprocessors, programmable logic units, dedicated logic and associated circuitry. This unit interfaces with the system or input/output bus 34 and, in one configuration, also interfaces with any disk caches 30 which may be present. The unit receives and processes requests and queries from the CPU 32 or network interface 36 and presents these requests in the appropriate form to approximate matching unit 24. Additionally the unit may aid in passing the results of the requests to either or both the disk cache 30 and/or the CPU 32 (by way of the bus 34).

More particularly, as shown in Figure 1A, a conventional rigid disk drive may have a plurality of rotating disks with multiple transducers accessing each disk. Each of these transducers typically has its output feeding analog signal electronics, such as amplifiers. This is represented at point A in the Figure. The Approximate Matching and Pre-fetch Processor (AMPP) 20, or the AMU 24, may be installed at point A for analog matching. As further shown in Figure 1A, typically the outputs of the analog circuitry are selectively provided to a single digital decoder which then processes one such output. This is represented at point B in the Figure. This digital output is typically then sent through error correction circuitry (ECC) and at its output C is then passed on to the bus 34 or disk cache 30. For purposes of the present invention, it may be desirable to provide multiple parallel paths for data by providing multiple digital decoders and ECC's. The Approximate Matching and Pre-fetch Processor (AMPP) 20, or the AMU 24, may be installed at either of points B or C for digital matching, and for multiple data output, multiple AMPP's or AMU's may also be provided.

Also shown in Figure 1 is a network interface 36 interconnecting the present invention 20, a network 38 which may be a LAN, WAN, Internet, etc. and to which other computer systems 40 may be connected. With this arrangement, other computer systems 40 may 5 conveniently also access the data stored on the mass storage medium 26 through the present invention 20. More specific examples are given below.

As shown in Figures 2-4, the present invention may be implemented in a variety of computer and network configurations. As 10 shown in Figure 2, the present invention may be provided as part of a stand-alone computer system 42 comprising a CPU 44 connected to a system bus 46 which then accesses a mass storage medium 48 having the invention as disclosed herein.

As shown in Figure 3, the mass storage medium 50 coupled with 15 the present invention may be itself connected directly to a network 52 over which a plurality of independent computers or CPU's 54 may then access the intelligent mass storage medium 50. The intelligent mass storage medium 50 may itself be comprised of a bank of hard disk drives comprising a RAID, disk farm, or some other massively parallel 20 memory device configuration to provide access and approximate matching capabilities to enormous amounts of data at significantly reduced access times.

As shown in Figure 4, an intelligent mass storage medium 56 equipped with the present invention may be connected to a network 58 25 as a Network Attached Storage Device (NASD) such that over the network 58 a plurality of stand-alone computers 60 may have access thereto. With the configuration as shown in Figure 4, it is contemplated that each data storage device, represented as a disk in 30 Figure 4, would be accessible from any processor connected to the network. One such configuration would include assigning a unique IP address or other network address to each data storage medium.

The configurations as exemplified by those shown in Figures 2-4 represent only examples of the various computer and network 35 configurations with which the present invention would be compatible and highly useful. Others would be apparent to those having skill in the art and the present invention is not intended to be limited through the examples as shown herein which are meant to be instead 40 illustrative of the versatility of the present invention.

As shown in Figure 5, the method of the present invention is 40 described alternatively with respect to whether an analog or digital key is used. However, beginning at the start of the method, a CPU

performs certain functions during which it may choose to access data stored in a mass storage medium. Typically, the CPU runs an application 62 which may be a DNA search, an Internet search, an analog voice search, a fingerprint search, an image search, or some 5 other such search during which an approximate (or exact) match to data is desired. In executing that higher order language application, the CPU would typically execute a request or query to obtain information from the storage device 26. The query contains directives specifying various parameters which the disk control unit 10 28 and the approximate matching and prefetch processor 20 must have to properly obtain requested information from the storage device 26. Examples of parameters include but are not limited to the following: the starting location for scanning the storage device; the final 15 location after which (if there is not match) scanning is terminated; the key to be used in the scanning; a specification of the approximate nature of the matching; what information should be returned when a match occurs, etc. The sort of information that can be returned includes the address of the information where the match was found, or a sector, record, portion of record or other data 20 aggregate which contains the matched information. The data aggregate may also be dynamically specified in that the data returned on a match may be specified to be between bounding data specifiers with the matched data contained within the bounding field. For example, looking for the word "terrorist" in a string of text might find the 25 approximate match, due to misspelling, of the word "terrerist", and return a data field which is defined by the surrounding sentence. Another query parameter would indicate whether the returned information should be sent to the system or input/output bus 34, or the disk cache 30.

30 The query will typically result in the execution of one or more operating system utilities. As an example of a higher level utility command, for the UNIX operating system, this could be modified versions of glimpse, find, grep, apropos. Other such functions causing the CPU to send commands 66 such as search, approximate 35 search, etc., to the approximate matching and pre-fetch processor of the present invention (see Fig. 1) with relevant portions of these commands also being sent to the disk controller 28 to, for example, initiate any mass storage media positioning activity 69 that is later required for properly reading information from the media.

40 At this point, depending upon the particular methodology desired to be implemented in the particular embodiment of the

invention, it would be necessary that an analog or digital key is determined. This key is data, exact or approximate, that corresponds to the data being searched for. For an analog key, this key may either be pre-stored such as in the mass storage medium, developed 5 using dedicated circuitry, or required to be generated. Should the analog key be pre-stored, a send pre-stored key step 68 would be performed by the microprocessor or programmable logic device 22 (see Fig. 1) which would transmit the key in digital and sampled format to the approximate matching unit 24 as shown in step 70. Alternatively, 10 should the analog key not be pre-stored, the analog key can be developed using one of a number of mechanisms, two of which are shown in Figure 5. In one the microprocessor 22 would write the key on the magnetic medium as at step 72 and then next read the key as at step 74 in order to generate an analog signal representation of the key. 15 In another as at step 71 the digital version of the key received from the CPU would be converted using appropriate digital to analog circuitry to an analog signal representation which would in turn be appropriately sampled. The key would then next be stored as a digital sample thereof as in step 70. Should a digital key be used, 20 it is only necessary that the microprocessor 22 store the digital key as at step 76 in approximate matching unit 24. It should be understood that depending upon the particular structures desired to be included for each of the microprocessor 22 and approximate matching unit 24, the key may reside in either or all of these 25 components, it merely being preferable to ultimately get the appropriate digital format for the key into the approximate matching unit 24 for comparison and correlation. Next, after the mass storage device reaches its starting location as at 79, the data stored on the mass storage medium 26 is continuously read as at step 78 to generate 30 a continuous stream signal representative of the data stored in the mass storage medium 26. Should an analog key have been used, this analog key may then be correlated with an analog read of data from the mass storage medium 26 as at step 80. While the inventors contemplate that any of many prior art comparators and correlation 35 circuitry could be used, for present purposes the inventors suggest that a digital sampling of the analog signal and key could be quite useful for performing such comparison and calculating the correlation coefficient, as explained below. It is noted that this analog signal generated from reading the data from mass storage medium 26 may be 40 conveniently generated by devices in the prior art from the reading of either analog or digital data, it not being necessary that a

digital key be used to match digital data as stored in mass storage medium 26. Alternatively, a correlation step 82 may be performed by matching the digital key with a stream of digital data as read from the mass storage medium 26. Note that the key reflects the inclusion 5 of approximate information. Thus, correlating this with information read from the storage medium enables approximate matching capabilities.

Decision logic 84 would next make an intelligent decision as to whether a portion of data approximately matches or does not 10 approximately match the key. Should a match be found, then the data is processed as at step 86 and the data requested by the query is sent to a disk cache 30, directly onto system bus 34, or otherwise be buffered or made available to a CPU 32, network interface 36, or otherwise as shown in Figures 1-4. A logical step 88 is suggested in 15 Figure 5 for returning to the continuous reading of data from the mass storage medium 26, indicating perhaps a "do" loop. However, it should be understood that this is a continuous process and that data is processed from the mass storage medium 26 as a stream and not in individualized chunks, frames, bytes, or other predetermined portions 20 of data. While this is not precluded, the present invention preferably allows a key to be in essence "slid" over a continuously varying data read signal such that there is no hesitation in reading data from the mass storage medium 26. There is no requirement to synchronize reading to the start or end of any multi-bit data 25 structure, or any other intermediate steps required to be performed as the data is compared continuously "on the fly" as it is read from the mass storage medium 26. This type of comparison and correlation may be referred to as a pattern match or comparison. Eventually, the data access is completed as at step 90 and the process completed.

30 The inventors herein have preliminarily tested the present invention and have generated preliminary data demonstrating its operability and effectiveness.

Figure 6 is a graphical representation of a measured analog signal output from a read/write head as the read/write head reads a 35 magnetic medium on which is stored a 10-bit digital key. As shown therein, there are peaks in an analog signal which, as known in the art, represents the true analog signal generated by a read/write head as data is read from a magnetic medium such as a hard disk. The scales shown in Figure 6 are volts along the vertical axis and tenths 40 of microseconds along the horizontal axis.

As shown in Figure 7, an analog signal is generated, again by a read/write head, as data is read from a pseudo-random binary sequence stored in a test portion of a magnetic medium. As can be seen in Figure 7, the read signal does not provide an ideal square wave 5 output when examined at this level.

Figure 8 is a graphical representation, with the horizontal scale expanded, to more specifically illustrate the overlap between approximately two bits of the 8-bit key and the corresponding two bits found in the pseudo-random binary sequence encoded at a 10 different location on the disk or magnetic medium.

Figure 9 is a graphical representation of a correlation coefficient calculated continuously as the comparison is made between the key and the continuous reading of data from the hard disk. This correlation coefficient is calculated by sampling the analog signals 15 at a high rate and using prior art signal processing correlation techniques. One such example may be found in Spatial Noise Phenomena of Longitudinal Magnetic Recording Media by Hoinville, Indeck and Muller, IEEE Transactions on Magnetics, Volume 28, no. 6, November 1992, the disclosure of which is incorporated herein by reference. A 20 prior example of a reading, comparison, and coefficient calculation method and apparatus may be found in one or more of one of the co-inventor's prior patents such as US Patent No. 5,740,244, the disclosure of which is incorporated herein by reference. The foregoing represent examples of devices and methods which may be used 25 to implement the present invention, however, as mentioned elsewhere herein, other similar devices and methods may be likewise used and the purposes of the invention fulfilled.

At approximately the point labeled 325, a distinct peak is noted which approaches 1, indicating a very close match between the 30 key and the pseudo-random binary sequence. Thus, the present invention has been demonstrated as capable of finding a key that is present in a pseudo-random digital sequence. Figure 9 is also 35 illustrative of the opportunity for approximate matching which is believed to be a powerful aspect of the present invention. Looking closely at Figure 9, it is noted that there are other lesser peaks that appear in the correlation coefficient. Thus, if a threshold of 0.4 were established as a decision point, then not only the peak occurring which approaches 1 would indicate a match or "hit" but also another five peaks would be indicative of a "hit". In this manner, a 40 desired coefficient value may be predetermined or adjusted as desired to suit particular search parameters. For example, when searching

for a particular word in a large body of text, lower correlation values may indicate the word is present but misspelled.

Figure 10 depicts the continuous calculation of a correlation coefficient between the same 8-bit key but with a different data set. 5 Again, a single match is picked up at approximately 200 microseconds. It is also noted that should a lower threshold be established additional hits would also be located in the pseudo-random data.

Various changes and modifications to the present invention would be apparent to those skilled in the art but yet which would not 10 depart from the spirit of the invention. Several of these alternatives have been indicated above. For example, all of the operations exemplified by the analog processing shown in Figures 6- 15 10, have their equivalent counterparts in the digital domain. Thus, approximate matching and correlation types of processing can be done on the standard digital representation of the analog bit patterns. This can also be achieved in a continuous fashion using tailored digital logic, microprocessors and digital signal processors, or alternative combinations. It is therefore the inventors' intention 20 that the present invention be limited solely by the scope of the claims appended hereto, and their legal equivalents.